

Amendments to the Specification

Please replace the paragraph that begins on Page 10, line 6 and carries over to Page 11, line 5

with the following marked-up replacement paragraph:

-- Measurement data is used to form cube structures within a spatially-enabled database system or systems, where these cube structures represent various aspects of the monitored business processes, according to preferred embodiments. These cube structures are also referred to herein as geospatial iceberg cubes or, equivalently, as geospatial cubes, where these terms may be defined as a graphical representation of data (e.g., service level management data) in multiple dimensions leveraging relational spatial structures. One or more reference objects is programmatically created, according to preferred embodiments, representing objectives for measurement data at a particular level of service. For example, a reference cube may be used that represents service objectives of a platinum service level, and another reference cube may be used to represent the gold service level. These reference objects are also referred to herein as “comparative objects” or [[a]] “boundary objects”, and may be structured as cubes or elements thereof (such as planes). Comparative objects may be redefined or modified programmatically (e.g., when SLA commitments are changed). In preferred embodiments, cubes containing measurement data are linked or interrelated to one another, thereby facilitating efficient drill-down; comparative objects may be linked or interrelated in this manner as well. In addition to, or instead of, building the comparative objects to represent objectives pertaining to SLAs, the comparative objects can be based on other factors such as feedback information from other systems (including historical data) and/or input from an administrator of the SLMS. (For example, the administrator might be queried for parameter values such as a desired upper and

lower bound on response time for each of the levels of service, and a reference object may then be constructed from these input values.) --

Please replace the paragraph on Page 12, lines 3 - 11 with the following marked-up replacement paragraph:

-- Fig. 1 provides a comparison of OLAP systems of the prior art and geospatial iceberg cubes of the present invention. As shown therein, an OLAP program is comprised of algorithms and data structures, whereas a program used with a geospatial iceberg cube is also comprised of algorithms and data structures but may now be an autonomic program (that is, a program that uses results of analyzing one or more cubes to autonomically modify system behavior). OLAP systems typically use matrices, records, and bitmaps for data storage, whereas geospatial iceberg cubes of the present invention use spatial cubes (i.e., [[cubes]] cube structures and operations available in prior art spatially-enabled database systems); data points, lines, and/or polygons; and possibly other spatial objects. --

Please replace the paragraph on Page 16, lines 3 - 18 with the following marked-up replacement paragraph:

-- Using the features of a spatially-enabled database system, for example, two planes can be compared to each other by invoking built-in functions. According to preferred embodiments of the present invention, planes are constructed from measurement data, and a collection of these planes forms a cube. Storing this data in a spatially-enabled database system (implementations of which are commercially available, as just discussed) provides a number of advantages. Existing

data normalization and data management features of the relational database technology in the spatially-enabled database management system can [[also]] be leveraged, thereby improving efficiency. The spatial extensions, geometric data types, grid indexing functions, user-defined functions, and/or built-in procedures of the database system can also be used. In this manner, operations on the stored data can use optimized built-in functions of the database system, rather than requiring an applications programmer to provide complex code in his/her application for interacting with the gathered measurement data used by preferred embodiments. As a result, programmer efficiency is increased and code complexity is reduced, thereby leading to decreased program development and support costs. (Furthermore, use of the optimized built-in database functions for interacting with the stored measurement data will typically increase the efficiency of other application programs and search utilities that may access this data.) --

Please replace the paragraph on Page 18, lines 2 - 7 with the following marked-up replacement paragraph:

-- Fig. 3 further illustrates that assembler “A1” provides materials to another assembler “A2”, and a probe “c” measures data pertaining to business processes between these entities. Yet another assembler [[“A4”]] “A3” receives materials from assembler “A2” and from supplier “S3”, and operations between these entities are measured by probes “d” and “e”, respectively. In this example, assembler “A3” distributes completed goods to three retailers “R1”, “R2”, and “R3”, and probes “f”, “g”, and “h” measure business processes between these entities. --

Please replace the paragraph that begins on Page 18, line 8 and carries over to Page 19, line 11

with the following marked-up replacement paragraph:

-- A number of terms used herein to describe preferred embodiments will now be defined in more detail.

- Key process indicator, or “KPI”: This term is used herein to generally signify a measurement that is representative of a business process, or a factor that is measured for a business process. The time required to complete an action may be a KPI in a time-sensitive business process. For example, with reference to Fig. 3, one KPI may be the time required by “S1” to deliver its raw materials to “A1” after “A1” initiates its order for those materials.
- Collaboration: Collaboration can be thought of as a “macro process” (i.e., a collection of one or more processes) composed of one or more KPIs. The time it takes for a particular collaboration to occur may be computed, in appropriate scenarios, as the sum of the KPIs which make up that collaboration. (There may also be scenarios in which the summation of all KPIs is not appropriate. For example, if the collaboration includes processes that operate concurrently, then another approach, such as adding the longest of the concurrent processes, may be preferable.) With reference to Fig. 3, for example, a collaboration may be defined as including delivery of materials from “S1” to “A1” and forwarding materials from “A1” to [[“A3”]] “A2”.
- Service Offering, or “SO”: A service offering is a group of one or more collaborations. Thus, providing completed goods to retailer “R1” in Fig. 3 may be considered as a service offering that includes all the processes whereby materials

are delivered to “A3” and the process whereby “A3” delivers the completed goods to “R1”. A particular collaboration may be a member of one or more service offerings. --

Please replace the paragraph on Page 21, lines 4 - 7 with the following marked-up replacement paragraph:

-- Fig. 4 shows a list 400 of the probes contained in the example environment in Fig. 3, a list 410 of KPIs that might be used in this environment, a list 420 of collaborations that might be used when evaluating business processes in this environment, and a list 430 of ~~services~~ service offerings that might be evaluated. (Note that lists 420 and 430 are not intended to be exhaustive.)

--

Please replace the paragraph that begins on Page 21, line 8 and carries over to Page 22, line 8 with the following marked-up replacement paragraph:

-- The list 400 of probes simply contains one entry for each probe. The list 410 of KPIs uses a notation “ $\Delta(\text{probe}_{\text{destination}} \text{probe}_{\text{source}})$ ” for representing a KPI that measures elapsed time, from a starting point prior to the link monitored by “ $\text{probe}_{\text{source}}$ ” through the link monitored by “ $\text{probe}_{\text{destination}}$ ”. Thus, for example, “ $\Delta(\text{ca})$ ” represents the time that elapses when supplier “S1” sends materials to assembler “A1” (as measured by probe “a”) and “A1” then forwards materials or goods to assembler “A2” (as measured by probe “c”). As a special case in this notation, “t” represents an initial time, and therefore the KPI “ $\Delta(\text{et})$ ” represents the time that elapses when supplier [[“S4”]] S3 sends materials to assembler “A3” (as measured by probe “e”). As another

example, the notation “ $\Delta(c-ab)$ ” might be used to represent only the time that elapses between a starting time when assembler “A1” receives materials from both S1 and S2 (as measured by probes “a” and “b”, respectively) and an ending time when materials or goods are then forwarded from “A1” to “A2”. In other words, only the elapsed time measured by probe “c”, on the link between “A1” and “A2”, is included for the KPI “ $\Delta(c-ab)$ ”. The number and type of KPIs to be created may vary widely among implementations of the present invention, and in list [[420]] 410, KPIs have not been constructed for every possible combination of flows. For example, it might be desirable to construct a KPI such as “ $\Delta(ag)$ ”, representing a duration of time from when supplier “S1” provides materials as input to the process, up through and including when retailer “R2” receives a product that has made its way through the system under evaluation. A KPI may be defined for any process of interest in a particular system, and alternative ways of defining and representing KPIs may be used in an embodiment of the present invention without deviating from the inventive techniques disclosed herein. --

Please replace the paragraph on Page 23, lines 2 - 13 with the following marked-up replacement paragraph:

-- In preferred embodiments, each KPI, each collaboration, and each service offering is associated with at least one threshold that is based on an SLA. Where multiple different levels of service are provided (e.g., platinum, gold, etc.), then each KPI, collaboration, and service offering is typically associated with a different threshold for each of these levels of service level. With reference to collaborations, for example, if the number of failures for a particular collaboration exceeds a certain value (where this value is normally specified in the SLA), then the SLA

commitment has not been met and the customer is typically due for some type of compensation. (Owing a reduced fee to the service provider is one type of compensation that might be applicable.) When a collaboration measurement pertains to a unit of time, then a collaboration failure may be a collaboration that fails to complete within the prescribed time limit. The thresholds for each KPI, each collaboration, and each service offering are preferably stored in the previously-discussed tables that were referred to above when discussing “collections 1 - 6”. --

Please replace the paragraph that begins on Page 26, line 9 and carries over to Page 27, line 4 with the following marked-up replacement paragraph:

-- Discussions herein are primarily in terms of ensuring that service is “good enough” to meet commitments. However, evaluations performed using embodiments of the present invention may be directed not only toward detecting failures to achieve SLA commitments, but also to areas where SLA commitments have been exceeded. That is, as a matter of good business practice, it may be desirable to not only avoid giving service that is “too bad” (meaning that the customer is unhappy and/or is entitled to a rebate or other compensation), but also to avoid giving service that [[it]] is “too good”. Service that is “too good” arises when a customer receives a level of service beyond what the customer paid for. For example, if the customer pays for silver level service but achieves results falling within the parameters of platinum level service, then the service provider is effectively “giving away” its higher-level service. In addition to lost revenue opportunities for the service provider, this situation may cause customers at the higher levels of service to become dissatisfied, even though their SLA commitments are met, if they learn that other customers also receive this higher level of service without paying its higher price. Thus, while discussions herein

are primarily in terms of evaluating when measurements do not rise to the level in an SLA commitment, it will be obvious in view of the teachings herein how embodiments may also, or alternatively, be used to evaluate whether measurements have been exceeded in the opposite sense. --

Please replace the paragraph on Page 28, lines 13 - 19 with the following marked-up replacement paragraph:

-- As with the cube in Fig. 5A, the cube in Fig. 6A can be used to determine failures at a given point in time, as well as failure trends over time, and the cube is preferably constructed from a collection of planes created over a relatively short time interval. The second-tier cube can then be compared to a second-tier reference cube and/or other reference objects. This second-tier reference cube is preferably constructed in a similar manner to the first-tier reference cube, except that the second-tier reference cube contains planes that graph the maximum number of allowable collaboration failures in each collaboration “A” through “Z” over the time interval T_0 through T_n .

--

Please replace the paragraph that begins on Page 29, line 11 and carries over to Page 30, line 10 with the following marked-up replacement paragraph:

-- In Fig. 6B, the second-tier cube from Fig. 6A is shown with a reference cube interposed. For purposes of illustration, dark shading has been used at reference number 610 in Fig. 6B to show how the reference cube intersects, or aligns with, the second-tier cube from Fig. 6A. Figs. 6C and 6D provide sample planes to further illustrate the sample reference cube. As

shown at reference number 620 in Fig. 6C, at some time T_j , a fairly high number of failures are tolerated: in this example, 6 failures are tolerated in each collaboration. (Sample values are provided on the y-axis in Figs. 6C and 6D merely for reference.) Fig. 6D represents a later time T_k , and as shown at reference number 630 in Fig. 6D, the number of failures that are tolerated at this later time is significantly less (2 for each collaboration, in this example) [[that]] than are tolerated at time T_j . It should be noted that while the sample reference planes shown in Figs. 6C and 6D have a horizontal line 620, 630 that illustrates uniform tolerance levels across all collaborations, this is merely an example. An actual reference cube may be comprised of planes where the tolerance varies from one collaboration to another. Thus, a line string or a collection of points may form an alternative plane for the (x,y) axes. Furthermore, more than one reference cube may be defined, where each such cube is constructed with reference to a different service offering. That is, one reference cube specifying tolerance values for platinum level customers might be interposed over the set of collaborations in the second-tier cube in Fig. 6A when drilling down from a platinum level service offering of the first-tier cube, and if a gold level service offering is also defined for those collaborations, then a different reference cube specifying different tolerance values may be interposed on the second-level cube of Fig. 6A when drilling down from the gold level in the first-tier cube. --

Please replace the paragraph that begins on Page 31, line 15 and carries over to Page 32, line 11 with the following marked-up replacement paragraph:

-- The third-tier cube in Fig. 7A can therefore be analyzed to determine deviations from objectives at a point in time as well as trends, over time, in these deviations. One of the (x,y)

planes, as illustrated by Fig. 7B, may be used for the point-in-time analysis of a collection of KPIs (as discussed above). Fig. 7C shows an alternative plane, as a vertical slice along the (y,z) axes of the third-tier cube, that may be used to evaluate trends for a selected KPI (i.e., a selected value from the x-axis) over time. Suppose this plane corresponds to KPI “A”. Fig. 7C shows that KPI “A” deviated, in this sample, from its objectives at each time interval T_0 through T_n . (Obviously, there may be cases where the objective was met, and in this case, a point may be placed directly on the z axis to indicate a (y,z) value of $(0, T_x)$.) Yet another alternative plane is illustrated in Fig. 7D, representing a horizontal slice along the (x,z) axes of the third-tier cube, for some particular deviation value (i.e., for some value from the y-axis). Suppose that this plane corresponds to .20 on the y-axis. The plane then has points for each of the KPIs that deviated from its objective by 20 percent over the time interval T_0 to T_n . For purposes of illustration, Fig. 7D shows that KPI “B” had this deviation at time T_4 and T_8 ; KPI “N” had this deviation at time T_3 ; and, by absence of a point from KPI “W”, indicates that KPI “W” did not deviate from its objective by 20 percent at any time during time interval T_0 to T_n . (Sample values are presented in the planes of Figs. 7B - 7D merely for illustrative purposes, and have been selected at random.) --

Please replace the paragraph on Page 32, lines 12 - 20 with the following marked-up replacement paragraph:

-- In Fig. 7E, a reference object 710 for a first service level and a reference object 720 for a second service level are shown as being mapped against the third-tier cube from Fig. 7A. Each reference object 710, 720 is typically associated with a different service level for which the per-KPI SLA commitments are to be evaluated. In this example, reference object 710 corresponds

generally to a range or band of deviations, across all KPIs and all time intervals, that extends from a 20 percent deviation at the lower bound to an 80 percent deviation at the upper bound. So, for example, this reference object 710 is concerned with determining whether any of the KPIs had deviations within this 20 percent to 80 percent range during the entire measurement interval T_0 to T_n . --

Please replace the paragraph that begins on Page 33, line 15 and carries over to Page 34, line 1 with the following marked-up replacement paragraph:

-- Reference object 720 corresponds, in the example of Fig. [[6F]] 7E, to a subset of the KPIs, measured across all time intervals, and corresponds generally to a range of deviations that extends from a 40 percent deviation at the lower bound to a 60 percent deviation at the upper bound. So, for example, this reference object 720 is concerned with determining whether KPIs “C” through “Q” had deviations within this 40 percent to 60 percent range during time interval T_0 to T_n . Many other reference objects might be used alternatively, and although the sample reference objects 710, 720 encompass the entire time interval T_0 to T_n , other reference objects might be directed toward a subset of this time interval. --

Please replace the paragraph that begins on Page 35, line 17 and carries over to Page 36, line 6 with the following marked-up replacement paragraph:

-- The building of the cubes may be triggered in several ways, without deviating from the scope of the present invention. As one example, an administrator might invoke an analysis function to check recent system performance and/or to predict future system performance, where

this analysis function builds and evaluates cubes according to preferred embodiments of the present invention. As another example, invocation of the cube-building function might be timer-driven, whereby the building is performed at regularly-scheduled points in time. Or, as yet another example, an event-driven approach might be used, whereby (for example) one or more threshold values are set, and the cube-building function is invoked when any of these thresholds is exceeded. Optionally, the scheduled times used in a timer-driven approach and the events and threshold values in [[this]] an event-driven approach may be configurable. --

Please replace the paragraph that begins on Page 38, line 15 and carries over to Page 39, line 8 with the following marked-up replacement paragraph:

-- Block 825 then populates a first-tier cube from the values obtained in Block 820, and in Block 830, this first-tier cube is displayed. One or more aspects of this first-tier cube are then evaluated, for example by a human user or a programmatic process. Based on this evaluation, in Block 835, a decision is made as to whether it is desirable to drill down for further detail -- i.e., to see the collaborations for a selected service [[level]] offering. For example, a selected one of the service offerings may have gone out of bounds of its SLA commitments (as may be determined with reference to SLA information stored in collection 6), and it may be desirable to drill down to the collaborations to see where problems may have arisen. Other heuristics may be evaluated as well. As one example, failure behavior at selected intervals may be evaluated for each of the service offerings, and if an undesirable trend is found, it may be desirable to drill down for further analysis. As another approach, each SO may be evaluated in turn by drilling down into its collaborations. If the test in Block 835 has a negative result, then Block 840 checks to see if

analysis of the system is finished. If so, then processing of Fig. 8 exits (Block 845); otherwise, control returns to Block 805 to continue sampling probes for the system. --

Please replace the paragraph on Page 42, lines 6 - 15 with the following marked-up replacement paragraph:

-- An embodiment of the present invention may alternatively or additionally allow for bottom-up processing instead of the drill-down processing that has been described thus far. So, for example, measurements for the KPIs may be created and summed/aggregated, and one or more third-tier cubes may be presented to represent this data. An upwards progression can then be made, if desired, to [[see]] view data pertaining to a collaboration including selected KPIs and/or a service offering that includes selected collaborations. In this bottom-up approach, it may be desirable to provide the user with a table or other legend that reflects the mapping to the next-higher level; enabling the user to select an entry from this table or legend will typically be easier than requiring the user to individually select an appropriate group of elements from the currently-viewed level. --